

IN THE SPECIFICATION:

Please replace the following paragraph beginning at page 8, line 13 through at page 10, line 15 as follows:

The routine begins at step 20 by deploying the ASR engine, for example, at a customer installation in the field. It is assumed that the ASR engine as initially deployed is imperfect and must be tuned to increase its accuracy level. At step 22, a test is run to determine whether the ASR engine has been set for automatic adaptation according to the invention. This may be achieved, for example, by extending the engine's application programming interface (API) to include a flag that may be set (e.g., by the application developer or the user) to begin the adaptation process. The flag may be set remotely after the engine has been deployed. If the engine has been set for automatic adaptation, the routine continues in step 24. At this step, a test is made to determine whether a given data-collection period (e.g., a 24 hour period) has elapsed. If not, the routine continues at step 26 to store the recognition results, along with the associated speech waveform samples. In particular, during this step, the recognition results obtained by processing a live input 38 are saved to a disk file. As will be seen, the recognition results may include the actual results (i.e., the hypothesized spoken utterances) generated by the ASR engine, together with information such as confidence levels, n_best hypotheses, and other data which might be used as input to the adaptation algorithms in step 30.

In a variant of step 26, the system may be configured to save a "downstream" version of the speech data (e.g., cepstral coefficients), instead of the "raw" digitized speech waveform samples 39. This is another advantage of the present invention. In particular, because there is no requirement for humans to listen to the speech data,

significant data-reduction may be obtained by storing only the form of the speech data that is required for executing the adaptation algorithms. This advantage can result in reduction in costs for computer equipment, including CPUs, IC memory, and hard disks.

Step 26 cycles until the result of the test at step 24 is positive. At this point, the routine continues at step 28 to retrieve the information saved during the time period. At step 30, an adaptation algorithm (or a plurality of adaptation algorithms) is executed against the information to increase the accuracy of the engine. As will be seen, this algorithm may be based on an acoustic model 33 (e.g., Hidden Markov Modeling 34), a language model (e.g., Word Bigram Statistics 35), a pronunciation model 36 (e.g., Phonetic Transcription 37), or some combination of these different model types. At step 32, the so-tuned recognition engine is then re-installed in the application, presumably with better accuracy and more efficient use of computing resources than the original engine.